

Bounded Distortion Mapping Spaces for Triangular Meshes

Yaron Lipman
Weizmann Institute of Science

Abstract

The problem of mapping triangular meshes into the plane is fundamental in geometric modeling, where planar deformations and surface parameterizations are two prominent examples. Current methods for triangular mesh mappings cannot, in general, control the worst case distortion of all triangles nor guarantee injectivity.

This paper introduces a constructive definition of generic convex spaces of piecewise linear mappings with guarantees on the maximal conformal distortion, as-well as local and global injectivity of their maps. It is shown how common geometric processing objective functionals can be restricted to these new spaces, rather than to the entire space of piecewise linear mappings, to provide a bounded distortion version of popular algorithms.

Keywords: triangular meshes, bounded distortion, conformal distortion, quasi-conformal, bijective mapping

Links: [DL](#) [PDF](#)

1 Introduction

Triangular meshes are prevalent in image and geometry processing and frequently used as piecewise linear representations of planar domains $\Omega \subset \mathbb{R}^2$ and surfaces $\mathcal{S} \subset \mathbb{R}^3$. A common task is to map or deform a triangular mesh into a different configuration in the plane $\Omega' \subset \mathbb{R}^2$ with as little distortion as possible; two central scenarios are: 1) planar mesh deformation $\Omega \rightarrow \Omega'$; and 2) surface mesh parametrization $\mathcal{S} \rightarrow \Omega'$.

Triangular meshes are mapped by Continuous Piecewise Linear (CPL) mappings \mathcal{F}^M constructed by assigning an affine map per triangle and making sure the affine maps are continuous across edges. A common distortion measure for a CPL map is conformal distortion, intuitively defined as the change in aspect ratio of every transformed triangle. Producing mappings of general triangular meshes with bounded *maximal* conformal distortion is important in geometric processing/modeling, physical simulations and numerical analysis. Simply put, approximation properties, stability, convergence and even visual quality of the maps are directly related to the *worst* distorted face(s).

Commonly, algorithms for surface parametrization and planar deformations optimize some average (e.g., L^p) distortion measure over the space of all CPL maps and therefore find it hard to avoid high conformal distortions on a subset of triangles. For example, Figure 1(b) shows a deformation of a (well-shaped) 2D mesh (shown in (a)) with discrete harmonic mapping [Pinkall and Polthier 1993; Eck et al. 1995]; the mapping suffers from high conformal distortion and fold-overs due to the complexity of the domain (non-convex, multiply-connected).

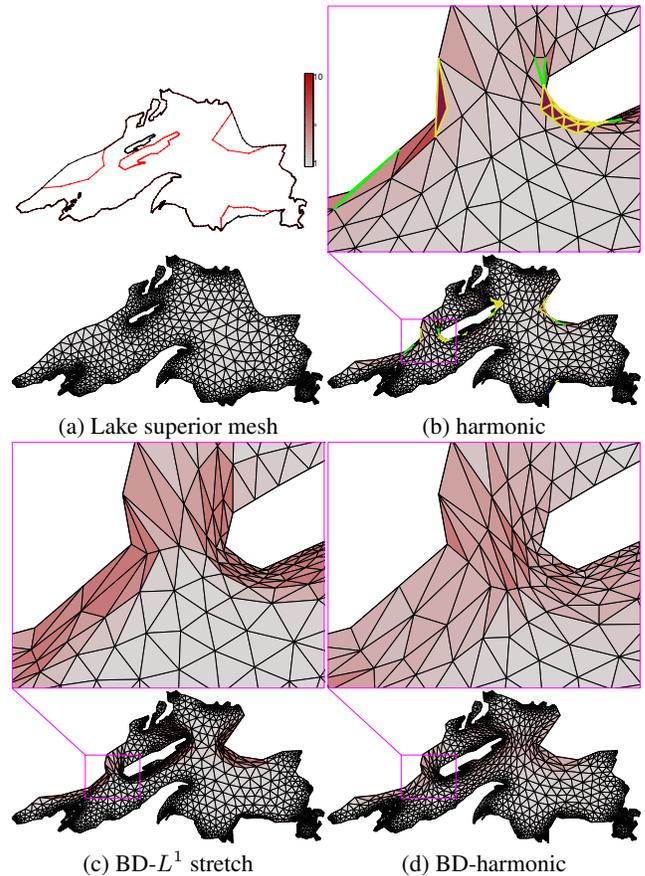


Figure 1: Deformation of the multiply-connected Superior Lake mesh by moving and scaling its center island and altering its shoreline, see (a) where the deformation constraints are in red. Using discrete harmonic maps, shown in (b), results in fold-overs and high distortion (redness indicates conformal distortion, flipped triangles are marked with yellow, highly-distorted ones in green). Optimizing the same energy (Dirichlet) over the bounded distortion mapping space is shown in (d), and produces a bounded conformal distortion bijection at the cost of only a slight increase (+4%) in the minimal Dirichlet energy. (c) shows the result of bounded distortion optimization of an alternative L^1 stretch functional.

The main goal of this paper is to devise a generic tool for constructing orientation preserving (i.e., no flips allowed) triangular mesh mappings into the plane with bounds on the worst-case conformal distortion. We denote this space of mappings $\mathcal{F}_C^M \subset \mathcal{F}^M$, where $C \geq 1$ is the desired upper bound on the maximal conformal distortion. Unfortunately, \mathcal{F}_C^M is a rather complicated non-convex space, and poses a challenge to work with. This paper suggests a simplified formulation of this space that allows easy characterization of convex subspaces $\mathcal{F}_C^{M, \boxplus} \subset \mathcal{F}_C^M$ that are conveniently formulated as a set of linear inequalities and convex quadratic cones. The subsets $\mathcal{F}_C^{M, \boxplus}$ have one degree of freedom per face, set by fixing a frame field (denoted by the symbol \boxplus) over the mesh. This frame field restricts the general space of bounded distortion maps \mathcal{F}_C^M by limiting the maximal rotation angle allowed. Nevertheless, as we

will demonstrate, it is still large enough to be of interest. For example, this space can be used in combination with known geometric processing functionals to build a “bounded distortion version” of existing algorithms; Figure 1(d), for example, shows the result of optimizing the Dirichlet energy (used to define the discrete harmonic mapping in (b)) over $\mathcal{F}_C^{\mathbf{M},\mathbb{B}}$. The resulting map has a uniform bound on the conformal distortion of all triangles and is guaranteed to be a bijection at the cost of increasing the Dirichlet energy by merely 4%. (c) shows the result of optimizing a different functional, namely L^1 stretch, over $\mathcal{F}_C^{\mathbf{M},\mathbb{B}}$. Although it has the same guarantees as (d), the L^1 nature of this functional tends to concentrate distortion and provides a different type of deformation.

Besides controlling the maximal conformal distortion, the mapping space $\mathcal{F}_C^{\mathbf{M},\mathbb{B}}$ has other applications in geometry processing: 1) it guarantees local and global bijectivity of CPL mappings (even for non-convex and multiply-connected domains); 2) it can be augmented to incorporate other distortion measures (e.g., area); 3) provide means to approximate extremal quasi-conformal maps (here we provide numerical evidence and partial theoretical analysis); and 4) can be used for 2D mesh improvement.

The paper is organized as-follows: we review the relevant previous work in Section 2, set some notation in Section 3, introduce the bounded distortion mapping spaces in Section 4, discuss their bijectivity properties in Section 5, review applications of the mapping spaces in Section 6, discuss limitations in Section 7, and conclude in Section 8.

2 Previous work

Low distortion planar deformations. Controlling distortion in planar deformation/warping is a long standing research problem. Smooth energies such as Biharmonic [Bookstein 1989; Jacobson et al. 2011] can construct intuitive smooth deformations but in general do not control distortion. As-Similar-As-Possible (ASAP) and As-Rigid-As-Possible (ARAP) [Alexa et al. 2000; Igarashi et al. 2005; Schaefer et al. 2006] are two paradigms that require each face of a mesh to be mapped, to the extent possible, by a similarity or rigid-motion (respectively). The ARAP model was elaborated in follow-up works [Sorkine and Alexa 2007; Chao et al. 2010; Solomon et al. 2011]. Although ASAP and ARAP algorithms can produce low distortion on average, no guarantee is provided on individual mesh faces. Conformal mappings [Lipman et al. 2008; Weber and Gotsman 2010] can guarantee zero conformal distortion at places the mapping is not singular, however conformal maps have a small number of degrees of freedom and cannot be both bijective and interpolate more than a handful of points (e.g., one to three in simply connected domains). Quasi-Conformal (QC) mappings are mappings of bounded conformal distortion and can be seen as a natural generalization of conformal mappings [Ahlfors 1966]. Due to their favorable properties, approximations of QC maps were developed in numerical analysis literature [Gaidashev and Khmelev 2008]. Recent works in graphics use them for 2D mappings [Zeng and Gu 2011; Lipman et al. 2012], however, either no explicit bound on conformal distortion of the faces can be given, or the problem is solved in a particular sub-class of mappings (e.g., for four interpolation points). Lastly, a somewhat related recent paper by Johnen *et al.* [2012] showed how to compute bounds on the jacobian of curvilinear finite element taking boundary straight-line elements to curved ones.

Surface parametrization. Due to its many applications in graphics and related fields, surface parametrization is one of the most researched geometry processing problems and, currently, a wide spectrum of elegant and powerful algorithms for triangular mesh parameterization is available [Floater and Hormann 2005; Sheffer

et al. 2006]. Conformal distortion is probably the most popular distortion measure, although other objectives, such as As-Rigid-As-Possible energy functionals have been suggested as well [Liu et al. 2008]. Linear conformal methods try to minimize square deviation from satisfying the Cauchy-Riemann equations [Levy et al. 2002; Desbrun et al. 2002]. Non-linear methods usually either directly optimize distortion functionals [Maillot et al. 1993; Sander et al. 2001; Hormann and Greiner 2000] or work in the space of admissible metrics (angles/edge lengths) and conformally flatten the surface [Sheffer et al. 2005; Kharevych et al. 2006; Jin et al. 2007; Ben-chen et al. 2008; Springborn et al. 2008]. However, no bound or control of the maximal conformal distortion is available. One exception is [Sorkine et al. 2002] where a greedy optimizer flattens faces until distortion threshold is met, however there is no control over the number of seams or their location on the surface.

Global bijective mappings of triangular meshes. Bijectively mapping a surface mesh or a planar mesh into another planar domain is in many senses an un-solved problem in geometry processing and modeling. In the case of a convex target domain, Floater [2003] showed a principled way to build such maps using convex combination mappings, generalizing the classical result by Tutte [1963]. Unfortunately, Floater’s construction does not extend to non-convex target domains. Gortler *et al.* [2006] studied the conditions under which convex combination mappings are injective when the target domain is non-convex. A different partial solution in case the source domain is planar as well, is compatible meshing of the two domains [Aronov et al. 1993]. For example, Alexa *et al.* [2000] constructed such compatible triangulation for defining a bijective mapping between two domains in the plane. However, these methods do not provide a guarantee regarding the distortion of the map computed, nor do they provide a collection of possible solutions to choose from (e.g., to find a “smooth” solution). Furthermore, the generalization to surface meshes embedded in \mathbb{R}^3 is not evident. In this paper we show that the mapping spaces $\mathcal{F}_C^{\mathbf{M},\mathbb{B}}$ can be used to tackle the global bijective mapping problem: we prove that as long as the boundary of the source domain is mapped bijectively to a simple¹ polygonal curve, then *any* map $\phi \in \mathcal{F}_C^{\mathbf{M},\mathbb{B}}$ is a global bijection between the source domain and the target domain bounded by this curve. A generalization to multiply-connected domains is also provided.

Meshing of planar domains is a well-studied problem in computational geometry, and algorithms with guarantees on the shape of the triangles are known [Bern and Eppstein 1995]. The problem of 2D meshing involves either Steiner point insertion [Ruppert 1995; Shewchuk 1996] or mesh improvement [Amenta et al. 1997]. We show that if the distortion measure of interest is the conformal distortion of a mesh triangle w.r.t. a perfect equilateral, then the mapping spaces $\mathcal{F}_C^{\mathbf{M},\mathbb{B}}$ can be used to improve meshes produced by current meshing algorithms.

3 Notation

A triangular mesh $\mathbf{M} = (\mathbf{V}, \mathbf{E}, \mathbf{F})$ is an oriented piecewise linear 2-manifold that consists of a set of vertices $\mathbf{V} = \{v_i\}$, edges $\mathbf{E} = \{e_\ell\}$ and oriented faces (triangles) $\mathbf{F} = \{f_j\}$. In this paper we will consider both meshes embedded in 3-space $\mathbf{V} \subset \mathbb{R}^3$ and planar meshes $\mathbf{V} \subset \mathbb{R}^2$. Since a triangular mesh is a piecewise linear object consisting of triangles, the most natural mapping space of a mesh is the (standard) Continuous Piecewise Linear (CPL) mappings, namely, each face f_j is mapped with an affine map \mathcal{A}_j and the different affine maps agree on common edges and thus construct a globally continuous map. We denote the collection of such mappings of \mathbf{M} into the plane by $\mathcal{F}^{\mathbf{M}}$.

¹Simple polygon is a closed polygonal curve without self-intersections.

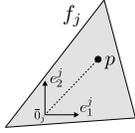
4 Bounded distortion mapping spaces

In this section we provide a simplified formulation to the bounded distortion mapping spaces of triangular meshes and extract maximal convex subsets (maximal in this formulation) that form the central tool of this paper. More specifically we: 1) study the space of affine transformations $\mathcal{A}_C(f_j)$ taking a single face $f_j \in \mathbf{F}$ into the plane without flipping it or conformally distorting it by more than a fixed constant $C \geq 1$; 2) provide a characterization of the maximal subspaces $\mathcal{A}_C^{\boxplus}(f_j) \subset \mathcal{A}_C(f_j)$ defined uniquely by setting a frame \boxplus_j in the face f_j ; 3) collect $\mathcal{A}_C^{\boxplus}(f_j)$ over all faces $f_j \in \mathbf{F}$, and incorporate continuity constraints to define maximal convex subspaces $\mathcal{F}_C^{\mathbf{M},\boxplus} \subset \mathcal{F}_C^{\mathbf{M}}$ of bounded distortion CPL mappings of the mesh \mathbf{M} ; and 4) discuss optimization over the mapping spaces $\mathcal{F}_C^{\mathbf{M},\boxplus}$.

Affine transformations with bounded conformal distortion.

We consider a single face $f_j \in \mathbf{F}$ and parameterize the possible affine maps $\mathcal{A}_j : f_j \rightarrow \mathbb{R}^2$ by selecting (arbitrarily for now) a local coordinate frame $\boxplus_j := [\bar{0}_j, \{e_1^j, e_2^j\}]$, where $\bar{0}_j$ is the origin and $\{e_1^j, e_2^j\}$ is a positively oriented orthonormal basis (see inset).

We write each point $p \in f_j$ in this frame $p = \bar{0}_j + e_1^j[p]_1 + e_2^j[p]_2$, where $[p] = ([p]_1, [p]_2)^t$ is the coordinate vector of p in the frame \boxplus_j . An affine map $\mathcal{A}_j : f_j \rightarrow \mathbb{R}^2$ can now be written in coordinates (using the standard basis in \mathbb{R}^2)



$$\mathcal{A}_j(p) = \begin{pmatrix} a_j & b_j \\ c_j & d_j \end{pmatrix} \begin{pmatrix} [p]_1 \\ [p]_2 \end{pmatrix} + \begin{pmatrix} t_{j,1} \\ t_{j,2} \end{pmatrix} = A_j[p] + T_j, \quad (4.1)$$

where $a_j, b_j, c_j, d_j, t_{j,1}, t_{j,2} \in \mathbb{R}$ and on the right hand side we have used matrix notation where $A_j = \begin{pmatrix} a_j & b_j \\ c_j & d_j \end{pmatrix}$, and $T_j = \begin{pmatrix} t_{j,1} \\ t_{j,2} \end{pmatrix}$.

We define the space $\mathcal{A}_C(f_j)$ to contain affine transforms \mathcal{A}_j that are orientation preserving (i.e., do not flip f_j) with bounded conformal distortion $C \geq 1$, that is we ask that

$$D(\mathcal{A}_j) \leq C \quad (4.2)$$

$$\det(\mathcal{A}_j) > 0, \quad (4.3)$$

where A_j is the linear component of \mathcal{A}_j , and $D(\mathcal{A}_j) = D(A_j) = \Gamma_j/\gamma_j$ is the conformal distortion of the affine map \mathcal{A}_j ($\Gamma_j \geq \gamma_j$ are the larger and smaller singular values of A_j). The conformal distortion $D(\mathcal{A}_j)$ measures the aspect ratio distortion caused by the affine map \mathcal{A}_j . Unfortunately, these conditions are non-linear and non-convex. Our general plan is however to characterize the largest effective convex piece of the space defined by these conditions. For understanding which part should we cut, we first need to transform the conditions (4.2),(4.3) into a ‘‘canonical’’ form. We do that next.

Taking the straightforward route and expressing conditions (4.2)-(4.3) in terms of a_j, b_j, c_j, d_j appearing in eq. (4.1) leads to cumbersome expressions that are hard to make sense of and/or work with. Instead we suggest to write \mathcal{A}_j in a different way, motivated by standard complex variables theory, that will considerably simplify the task of bounding the conformal distortion. Let us write

$$A_j = \begin{pmatrix} a_j + c_j & d_j - b_j \\ d_j + b_j & a_j - c_j \end{pmatrix}. \quad (4.4)$$

Note that we still use four parameters $a_j, b_j, c_j, d_j \in \mathbb{R}$ to represent A_j , and there is a unique way to write each real 2×2 matrix in this way. The affine map \mathcal{A}_j can now be written as

$$\mathcal{A}_j([p]) = \begin{pmatrix} a_j & -b_j \\ b_j & a_j \end{pmatrix} \begin{pmatrix} [p]_1 \\ [p]_2 \end{pmatrix} + \begin{pmatrix} c_j & d_j \\ d_j & -c_j \end{pmatrix} \begin{pmatrix} [p]_1 \\ [p]_2 \end{pmatrix} + \begin{pmatrix} t_{j,1} \\ t_{j,2} \end{pmatrix}, \quad (4.5)$$

where the matrix $\begin{pmatrix} a_j & -b_j \\ b_j & a_j \end{pmatrix}$ is the similarity part of A_j and $\begin{pmatrix} c_j & d_j \\ d_j & -c_j \end{pmatrix}$ the anti-similarity part. In fact, the similarity part is the *closest* similarity to A_j in the Frobenious norm, and a similar statement holds for the anti-similarity part. As it happens, A_j can be written as the sum of its closest similarity and anti-similarity transformations. This leads to particularly simple expressions for the singular values of A_j :

$$\Gamma_j = \sqrt{a_j^2 + b_j^2} + \sqrt{c_j^2 + d_j^2},$$

and

$$\gamma_j = \left| \sqrt{a_j^2 + b_j^2} - \sqrt{c_j^2 + d_j^2} \right|.$$

It is more convenient to write these equations with complex numbers; we set $\alpha_j = a_j + \mathbf{i}b_j$, and $\beta_j = c_j + \mathbf{i}d_j$, $\delta_j = t_{j,1} + \mathbf{i}t_{j,2}$, and write the coordinates of p as a complex number $[p] = [p]_1 + \mathbf{i}[p]_2$, this leads to an equivalent formulation of (4.5):

$$\mathcal{A}_j([p]) = \alpha_j [p] + \beta_j \bar{[p]} + \delta_j. \quad (4.6)$$

The singular values are now

$$\Gamma_j = |\alpha_j| + |\beta_j| \quad (4.7)$$

$$\gamma_j = \left| |\alpha_j| - |\beta_j| \right|, \quad (4.8)$$

and the conformal distortion of \mathcal{A}_j :

$$D(\mathcal{A}_j) = D(A_j) = \frac{\Gamma_j}{\gamma_j} = \frac{|\alpha_j| + |\beta_j|}{\left| |\alpha_j| - |\beta_j| \right|}, \quad (4.9)$$

where $|\alpha_j| = \sqrt{a_j^2 + b_j^2}$ and $|\beta_j| = \sqrt{c_j^2 + d_j^2}$. The condition that \mathcal{A}_j is orientation preserving is: $\det A_j = a_j^2 - c_j^2 - (d_j^2 - b_j^2) = |\alpha_j|^2 - |\beta_j|^2 > 0$. Conditions (4.2)-(4.3) are now equivalent to:

$$|\beta_j| \leq \frac{C-1}{C+1} |\alpha_j| \quad (4.10)$$

$$|\beta_j| < |\alpha_j|, \quad (4.11)$$

where $C \geq 1$ is the conformal distortion bound. Introducing a new variable $r_j \in \mathbb{R}$ we can further simplify these conditions by separating α_j and β_j :

$$|\beta_j| \leq r_j \frac{C-1}{C+1} \quad (4.12)$$

$$|\alpha_j| \geq r_j \quad (4.13)$$

$$r_j > 0. \quad (4.14)$$

This means that the set of triplets $(\alpha_j, \beta_j, r_j) \in \mathbb{C} \times \mathbb{C} \times \mathbb{R}$ that satisfies the above inequalities parameterize the entire space $\mathcal{A}_C(f_j)$ of orientation preserving affine maps with bounded conformal distortion C . The benefit in the above representation is: condition (4.12) is a convex cone, while condition (4.13) is the complement of an open convex cone. In a sense, the non-convex bounded conformal distortion and orientation preserving conditions (4.2)-(4.3) are decomposed into their simplest convex (4.12), and non-convex (4.13) parts.

We would now like to cut the *maximal* convex chunk out of the non-convex space of affine maps defined by eqs. (4.12),(4.13),(4.14). The problem is now reduced to dealing with a single and rather simple non-convex condition (4.13): the largest convex piece we can carve out of a domain that is the complement of a cone is (any) half space supported on one of the cone’s rays. For reasons that will be clear soon we make a particular choice that leads to the

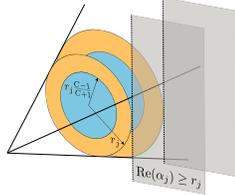
following definition of the sub-space $\mathcal{A}_C^{\boxplus j}(f_j) \subset \mathcal{A}_C(f_j)$:

$$|\beta_j| \leq r_j \frac{C-1}{C+1} \quad (4.15)$$

$$\operatorname{Re}(\alpha_j) \geq r_j \quad (4.16)$$

$$r_j > 0, \quad (4.17)$$

where $\operatorname{Re}(\alpha_j) = a_j$ is the real part of α_j . The inset provides a visual interpretation of the above inequalities: condition (4.15) requires β_j to lay in the blue cone (depicted with blue disk slices), while condition (4.16) asks α_j to lay in the gray half-space (illustrated with gray half-plane slices). For comparison, note that the ‘‘full’’ bounded conformal distortion conditions are different in that eq. (4.13) requires α_j to lay outside the yellow cone (depicted with yellow disks) which is more general than (4.16), however unfortunately non-convex.



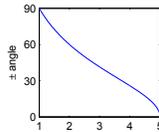
$\mathcal{A}_C^{\boxplus j}(f_j)$ is by construction convex and contains bounded distortion, orientation preserving affine transformations. However, a natural question is: *what mappings did we leave out when replacing $\mathcal{A}_C(f_j)$ with $\mathcal{A}_C^{\boxplus j}(f_j)$? and, why did we choose this particular half-space?* We address these questions next.

Properties of $\mathcal{A}_C^{\boxplus j}(f_j)$. The space $\mathcal{A}_C^{\boxplus j}(f_j)$ contains only a part of the full space $\mathcal{A}_C(f_j)$ of bounded distortion affine transforms.

In order to understand what exactly we included in $\mathcal{A}_C^{\boxplus j}(f_j)$ we provide the following proposition (proved in Appendix B):

Proposition 4.1. *Set $C \geq 1$. Then conditions (4.15), (4.16), (4.17) characterize the collection of orientation preserving affine transformations $\mathcal{A}_j(z) = \alpha_j z + \beta_j \bar{z} + \delta_j$ with conformal distortion $c \leq C$ under the following restriction on the similarity component α_j : $|\arg(\alpha_j)| \leq \cos^{-1}\left(\frac{C+1}{C-1} \frac{c-1}{c+1}\right)$. (The anti-similarity part β_j and the translational part δ_j are not restricted in any way w.r.t. the full space of bounded distortion affine space $\mathcal{A}_C(f_j)$.)*

To clarify the meaning of this proposition, let us note that the affine map \mathcal{A}_j is applied to the *coordinate vectors* of the vertices of face f_j w.r.t. its local frame \boxplus_j . That is, if we denote by j_ℓ , $\ell = 1, 2, 3$ the (ordered) indices of the vertices of the face f_j , then $(v_{j_1}, v_{j_2}, v_{j_3})$ are the vertices of face f_j , and the affine map \mathcal{A}_j is applied to the planar triangle $\Delta([v_{j_1}], [v_{j_2}], [v_{j_3}]) \subset \mathbb{R}^2$, where $[v_{j_\ell}] = [v_{j_\ell}]_1 + \mathbf{i}[v_{j_\ell}]_2$, $\ell = 1, 2, 3$ are the coordinates of the vertices in the frame \boxplus_j . Proposition 4.1 provides a restriction on the argument of the similarity component of the affine map, $\arg(\alpha_j)$, which is roughly the rotation angle of the affine map. For example, if we consider a conformal distortion bound of $C = 5$, then the space $\mathcal{A}_C^{\boxplus j}(f_j)$ contains affine transforms \mathcal{A}_j with rotations in the range depicted in the inset, where c denotes the conformal distortion of \mathcal{A}_j . In particular, an affine map \mathcal{A}_j with conformal distortion $c = 2$ can have a rotational angle $\arg(\alpha_j) \in [-60^\circ, 60^\circ]$. Perfect similarities can have rotation angles in $[-90^\circ, 90^\circ]$. It is clear now why we picked condition (4.16): it is a maximal convex piece of (4.13) that contains *symmetric* rotations around the identity.



Setting $\mathcal{F}_C^{\mathbf{M}, \boxplus}$. So far we have defined a convex space of affine transformations $\mathcal{A}_C^{\boxplus j}(f_j)$ per face $f_j \in \mathbf{M}$ via eqs. (4.15)-(4.17). The next step is to take the union of all these spaces over a mesh $\cup_{f_j \in \mathbf{F}} \mathcal{A}_C^{\boxplus j}(f_j)$ and construct $\mathcal{F}_C^{\mathbf{M}, \boxplus} \subset \mathcal{F}^{\mathbf{M}}$. The missing ingredient

is having the affine maps continuously agree across edges to achieve continuity. We take care of that while setting up $\mathcal{F}_C^{\mathbf{M}, \boxplus}$ next.

The variables we use to define $\phi \in \mathcal{F}_C^{\mathbf{M}, \boxplus}$ divide into two groups: 1) *face* variables, which include an affine map per triangle, are $\{\alpha_j, \beta_j, \delta_j, r_j\}_{j=1}^{|\mathbf{F}|}$ where $\alpha_j, \beta_j, \delta_j \in \mathbb{C}$, $r_j \in \mathbb{R}$ and $|\mathbf{F}|$ is the number of faces; and 2) *vertex* variables, which define the mapped location in the plane of each vertex $\phi(v_i) = u_i$, are $\{u_i\}_{i=1}^{|\mathbf{V}|}$, where $u_i \in \mathbb{C}$, and $|\mathbf{V}|$ is the number of vertices.

We assume to have a frame field $\boxplus = \{\boxplus_j\}_{j=1}^{|\mathbf{F}|}$ containing a frame per face. (We soon explain how to set the frame field \boxplus .) Then, to define $\mathcal{F}_C^{\mathbf{M}, \boxplus}$ we take one copy of equations (4.15)-(4.17) for every $f_j \in \mathbf{F}$, and add the following set of homogeneous linear equations forcing continuity:

$$\mathcal{A}_j([v_{j_\ell}]) = u_{j_\ell}, \quad f_j \in \mathbf{F}, \quad \ell = 1, 2, 3, \quad (4.18)$$

where \mathcal{A}_j is the affine map of face f_j defined in eq. (4.6). This is a set of $3|\mathbf{F}|$ linear (complex) equations in the variables $\{\alpha_j, \beta_j, \delta_j, r_j, u_i\}$ and hence convex. Adding it to the previous set of equations (4.15)-(4.17) maintain the convexity of the problem. Geometrically, one can imagine these linear hyper-spaces cutting through the different bounded distortion cones described before.

Our full representation is not very compact, but for certain cases (described later) we can make it more efficient by making the (standard) observation that $\{u_i\}$ sets all the degrees of freedom of CPL map of a mesh and the variables $\alpha_j, \beta_j, \delta_j$ can be expressed uniquely as linear combinations of u_{j_ℓ} to reduce the number of variables. Indeed, eq. (4.18) implies that for each face f_j , we can write $\alpha_j, \beta_j, \delta_j$ as linear combinations of u_{j_ℓ} , $j = 1, \dots, |\mathbf{F}|, \ell = 1, 2, 3$:

$$\begin{pmatrix} \alpha_j \\ \beta_j \\ \delta_j \end{pmatrix} = \begin{pmatrix} [v_{j_1}] & \overline{[v_{j_1}]} & 1 \\ [v_{j_2}] & \overline{[v_{j_2}]} & 1 \\ [v_{j_3}] & \overline{[v_{j_3}]} & 1 \end{pmatrix}^{-1} \begin{pmatrix} u_{j_1} \\ u_{j_2} \\ u_{j_3} \end{pmatrix}, \quad (4.19)$$

where $[v_{j_\ell}]$ are the complex coordinates of vertices j_ℓ in the local frame.

Adding positional constraints to the mapping space $\mathcal{F}_C^{\mathbf{M}, \boxplus}$ is instrumental in making it useful for practical mapping problems. Let us define the positional constraints on a subset of vertices $\mathbf{V}_{\text{con}} \subset \mathbf{V}$:

$$u_i = \hat{u}_i, \quad v_i \in \mathbf{V}_{\text{con}}, \quad (4.20)$$

where \hat{u}_i prescribe the constrained positions (anchors) in the plane. Although we will not use it in this paper, we note that adding rotational and/or scaling constraints for faces is also easy by prescribed the values $\alpha_j = \hat{\alpha}_j$ ($|\hat{\alpha}_j|$ is the scaling, and $\arg(\alpha_j)$ the rotation) for subset of faces $\mathbf{F}_{\text{con}} \subset \mathbf{F}$.

Framing of \mathbf{M} . To define the mapping space $\mathcal{F}_C^{\mathbf{M}, \boxplus}$, we need to set a frame field $\boxplus = \{\boxplus_j\}$ over the mesh \mathbf{M} . Every local frame \boxplus_j defines an initial rigid embedding of the face $f_j \in \mathbf{F}$ in the plane. Any map $\phi \in \mathcal{F}_C^{\mathbf{M}, \boxplus}$ is defined over these initial embeddings, mapping each embedded triangle $\Delta([v_{j_1}], [v_{j_2}], [v_{j_3}])$ to its final position while ‘‘stitching’’ the triangulation consistently and maintaining bounded conformal distortion C .

In the algorithms we develop in this paper we set the initial frame field $\boxplus = \{\boxplus_j\}$ by one of three possible ways: 1) if \mathbf{M} is planar, we can take constant frames (e.g., $\{1, \mathbf{i}\}$); 2) we can build a frame field based on a gradient field of some potential scalar function; or 3) we can choose a frame field based on an initial CPL map $\phi^* : \mathbf{M} \rightarrow \mathbb{R}^2$ or a set of provided affine maps $\mathcal{A}_j^* : f_j \rightarrow \mathbb{R}^2$ per face.

Let us elaborate on the third option. In this case we have an initial affine map \mathcal{A}_j^* per face f_j (we do not care if it is part of a con-

sistent CPL mapping ϕ^* or not). Proposition 4.1 indicates that the only constrained parts of the affine maps \mathcal{A}_j in the space $\mathcal{F}_C^{\mathbf{M},\mathbb{B}}$ (compared to the full space of affine transforms with bounded conformal distortion) are the arguments of the similarity components α_j (rotation angles). Furthermore, in order to get maximal “flexibility” over this initial transform we want to choose the frames \mathbb{B}_j such that writing \mathcal{A}_j^* w.r.t. these frames will result in $\alpha_j \in \mathbb{R}^+$ (that is $\arg(\alpha_j) = 0$). Practically, given an affine map $\mathcal{A}_j^*([p]) = \alpha_j^*[p] + \beta_j^*[p] + \delta_j^*$ w.r.t. the frame $\mathbb{B}_j^* = [\bar{0}_j^*, \{e_1^{*j}, e_2^{*j}\}]$ in face f_j , we rotate the basis $\{e_1^{*j}, e_2^{*j}\}$ by $\theta = -\arg(\alpha_j^*/|\alpha_j^*|)$ and set the new frame of face f_j to be

$$\mathbb{B}_j = e^{i\theta}\mathbb{B}_j^* = [\bar{0}_j^*, \{e^{i\theta}e_1^{*j}, e^{i\theta}e_2^{*j}\}]. \quad (4.21)$$

Optimization and implementation details. Equations (4.15) are quadratic convex cones and can be directly encoded in Second-Order Cone Programming (SOCP) optimizers such as MOSEK [Andersen and Andersen 1999]. Equations (4.16),(4.17),(4.18),(4.20) are all linear equalities and inequalities that are standard to encode. Later, in Section 6 we examine different functionals and optimize over $\mathcal{F}_C^{\mathbf{M},\mathbb{B}}$; SOCP optimizers can efficiently deal with linear objective functions which makes them suitable to optimize L^1 and L^∞ functionals.

For efficient optimization of L^2 functionals, it is useful to introduce another convex space of CPL maps $\mathcal{F}_C^{\mathbf{M},\mathbb{B},\infty} \subset \mathcal{F}_C^{\mathbf{M},\mathbb{B}}$ that is defined by replacing the quadratic cones with the L^∞ cones:

$$|\beta_j|_\infty \leq \frac{r_j}{\sqrt{2}} \frac{C-1}{C+1}, \quad (4.22)$$

where $|\beta_j|_\infty = \max\{|c_j|, |d_j|\}$. Optimizing an L^2 functional over $\mathcal{F}_C^{\mathbf{M},\mathbb{B},\infty}$ is a standard convex Quadratic Programming (QP) (with linear equalities and inequalities) which can be solved efficiently using interior point methods. We have used Matlab’s interior point algorithm for optimizing these problems. In this case it is also possible to reduce the number of variables in the optimization by using eq. (4.19), keeping only $\{u_i\}_{i=1}^{|\mathbf{M}|}$ as variables.

When using the more restricted space $\mathcal{F}_C^{\mathbf{M},\mathbb{B},\infty}$ we maintain our conformal distortion and orientation preservation guarantees, however we have a somewhat smaller space to work with. The tradeoff is explained in the following proposition proved in Appendix B:

Proposition 4.2. *Using the L^∞ cone condition (4.22) instead of the quadratic cone (4.15) defines a mapping space $\mathcal{F}_C^{\mathbf{M},\mathbb{B},\infty}$ that satisfies $\mathcal{F}_{C'}^{\mathbf{M},\mathbb{B}} \subset \mathcal{F}_C^{\mathbf{M},\mathbb{B},\infty} \subset \mathcal{F}_C^{\mathbf{M},\mathbb{B}}$, where $C' = \frac{(\sqrt{2}+1)C+(\sqrt{2}-1)}{(\sqrt{2}-1)C+(\sqrt{2}+1)}$.*

Simply put, $\mathcal{F}_{C'}^{\mathbf{M},\mathbb{B},\infty}$ is contained in $\mathcal{F}_C^{\mathbf{M},\mathbb{B}}$ and contains some smaller space $\mathcal{F}_{C'}^{\mathbf{M},\mathbb{B}}$, $C' < C$.

As final implementation detail, let us note that eq. (4.17) is numerically set by the linear inequality $r_j \geq \sigma$, where $\sigma > 0$ is some constant. In our implementation we chose it to be $\sigma = 10^{-8}$, however if no positional constraints are provided, like is common for free-boundary surface parametrization, then σ controls the global scale (note that without positional constraints the equations defining $\mathcal{F}_C^{\mathbf{M},\mathbb{B}}$, $\mathcal{F}_C^{\mathbf{M},\mathbb{B},\infty}$ are scale invariant) and therefore we set it to $\sigma = 1$ in these cases.

Resetting the frame-field and iterating. After optimization, where a solution $\phi \in \mathcal{F}_C^{\mathbf{M},\mathbb{B}}$ is found, eq. (4.21) can be used to reset the frames \mathbb{B}' and re-optimize over the new space $\mathcal{F}_C^{\mathbf{M},\mathbb{B}'}$. This will allow rotations that were not allowed in the former space. This procedure can be iterated until convergence. Our experience indicates that when using QP formulations, even when no feasible solution $\phi \in \mathcal{F}_C^{\mathbf{M},\mathbb{B}}$ is found with the current frame-field, it is often the case

that resetting the frame field using the intermediate output $\tilde{\phi}$ from the optimizer and resolving converges to finding a feasible solution. If that is not the case, increasing the conformal distortion bound C can also help to achieve a feasible solution. The reason is that increasing the conformal distortion bound C allows a wider rotation range in the convex subspace $\mathcal{F}_C^{\mathbf{M},\mathbb{B}}$ as Proposition 4.1 indicates. In practice, we performed such iterations until encountering one or two feasible iterations or alternatively exceeding some maximal number of iterations. In Section 6 we discuss and demonstrate the dependence of the final map ϕ on the initial frame field.

5 Bijectivity

Besides providing guarantees regarding conformal distortion and orientation preservation, the convex mapping spaces $\mathcal{F}_C^{\mathbf{M},\mathbb{B},\infty}, \mathcal{F}_C^{\mathbf{M},\mathbb{B}}$ can help us address an important problem of geometry processing, namely building globally bijective mappings. The results of this section actually apply to the bigger space of $\mathcal{F}_C^{\mathbf{M}}$ for all $C \geq 1$, that is to $\cup_{C \geq 1} \mathcal{F}_C^{\mathbf{M}}$, or differently put, to all orientation preserving CPL mappings (satisfying (4.3)). Let us denote this space by $\mathcal{F}_+^{\mathbf{M}}$. Since, in particular $\mathcal{F}_C^{\mathbf{M},\mathbb{B},\infty} \subset \mathcal{F}_C^{\mathbf{M},\mathbb{B}} \subset \mathcal{F}_+^{\mathbf{M}}$ for all $C \geq 1$, all subsequent results proved for $\mathcal{F}_+^{\mathbf{M}}$ are also valid for its convex subsets $\mathcal{F}_C^{\mathbf{M},\mathbb{B}}, \mathcal{F}_C^{\mathbf{M},\mathbb{B},\infty}$ and does not depend on any particular choice of frames \mathbb{B} .

We start with simply connected domains and later move to multiply-connected ones. Given a triangular mesh $\mathbf{M} \subset \mathbb{R}^d$, $d = 2, 3$ that is topologically a disk and a one-to-one mapping of its boundary polygon $\partial\mathbf{M}$ to a simple polygonal line P in the plane, we would like to construct an injective map that takes the interior of \mathbf{M} onto the domain Ω , defined as the domain bounded by P .

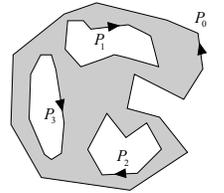
The orientation preservation condition enforced on all faces for mappings in $\mathcal{F}_+^{\mathbf{M}}$ ensures that on the *local* face level we are away from degeneracy and injective. However, with some effort, one can prove a much stronger global bijection result, in case the boundary conditions permit:

Theorem 5.1. *An orientation preserving CPL mapping (i.e., in $\mathcal{F}_+^{\mathbf{M}}$) that maps the boundary of \mathbf{M} bijectively to the boundary of Ω is a bijection from \mathbf{M} to Ω .*

Two comments are in order. First, note that we do not require the target domain Ω to be convex. Second, the positional constraints’ set can possibly contain extra conditions in interior points (in addition to the boundary constraints). This makes this method useful for constructions of bijective deformations required in practical settings like image deformation.

Theorem 5.1 can be generalized to multiply-connected domains:

Theorem 5.2. *Consider a multiply-connected mesh \mathbf{M} , with boundary $\partial\mathbf{M} = \cup_{\ell=0}^L P_\ell$ (P_0 is the outer boundary, and P_1, \dots, P_L the simple polygonal boundary of the holes, correctly oriented w.r.t. \mathbf{M} , see inset), and a homeomorphic, multiply-connected target domain $\Omega \subset \mathbb{R}^2$, with boundary $\partial\Omega = \cup_{\ell=0}^L P'_\ell$. An orientation preserving CPL mapping (i.e., in $\mathcal{F}_+^{\mathbf{M}}$) that maps the boundary of \mathbf{M} bijectively to the boundary of Ω is a bijection from \mathbf{M} to Ω .*



The proof requires several steps and we provide it in full in Appendix A. Figures 1,10 show examples of mapping multiply-connected domains bijectively using $\mathcal{F}_C^{\mathbf{M},\mathbb{B}}$ (or $\mathcal{F}_C^{\mathbf{M},\mathbb{B},\infty}$). We also show that deforming with discrete harmonic mappings creates foldovers and extreme conformal distortion (even on these well-shaped meshes).

6 Applications

We explored several applications of the bounded distortion mapping spaces $\mathcal{F}_C^{\mathbf{M},\mathbb{M}}, \mathcal{F}_C^{\mathbf{M},\mathbb{M},\infty}$ for triangular mesh mapping and related problems. A method for bounded distortion mesh mapping is defined by:

1. A bounded distortion mapping space $\mathcal{F}_C^{\mathbf{M},\mathbb{M}}$ (or $\mathcal{F}_C^{\mathbf{M},\mathbb{M},\infty}$) defined by setting a conformal distortion bound $C \geq 1$, and a frame field $\mathbb{M} = \{\mathbb{M}_j\}$.
2. A functional $E : \mathcal{F}^{\mathbf{M}} \rightarrow \mathbb{R}^+$ that provides a score to PLC maps.
3. Computing $\operatorname{argmin}_{\mathcal{F}_C^{\mathbf{M},\mathbb{M}}} E$ (or over $\mathcal{F}_C^{\mathbf{M},\mathbb{M},\infty}$).
4. (Optional) resetting the frame field (via eq. (4.21)) and iterating.

In this section we instantiate this framework to build specific algorithms, analyze them, and compare to the state of the art.

6.1 Planar deformations

In the first application we build algorithms for bounded distortion planar deformations. Our setting includes a mesh \mathbf{M} of a planar domain $\Omega \subset \mathbb{R}^2$, and a set of positional constraints. Our goal is to produce a bounded distortion deformation satisfying the constraints.

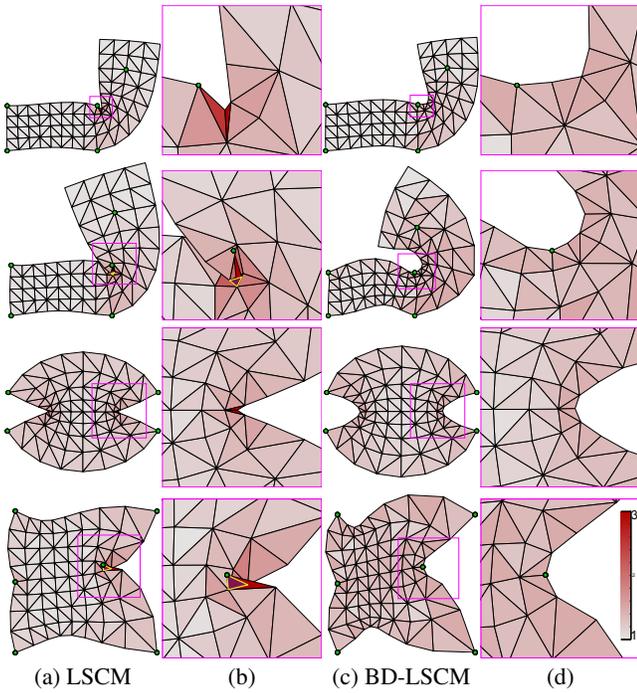


Figure 2: Deformation of regular grids with Least-Squares Conformal Maps (LSCM) [Levy et al. 2002] (a),(b); and Bounded Distortion LSCM (BD-LSCM) (c),(d).

As-similar-as-possible (ASAP) planar deformations. Levy et al. [2002] defined the Least Squares Conformal Maps (LSCM) on triangular meshes via a simple and elegant quadratic functional that measures the mean-square deviation from satisfying the Cauchy-Riemann (CR) equations. In our notation, the CR equations on every face are $\beta_j = 0$ (since $\partial_{\bar{z}}\mathcal{A}_j = \beta_j$ on face f_j), and Levy integrates their squared norms $|\beta_j|^2$, namely $E_{\text{LSCM}}(\phi) = \sum_j |\beta_j|^2 \text{Area}(f_j)$. However, small $|\beta_j|$ does not

mean that the conformal distortion (of the corresponding face) is low as $|\alpha_j|$ can be small as-well; as eq. (4.9) implies, small conformal distortion means that β has to be small w.r.t. α . For that reason LSCM can introduce arbitrary large conformal distortion and flipped triangles as Figure 2 (a),(b) demonstrate. It is possible to make sure the conformal distortion of LSCM is bounded and that no triangles are flipped by optimizing the LSCM energy over $\mathcal{F}_C^{\mathbf{M},\mathbb{M},\infty}$ rather than the general space of CPL maps $\mathcal{F}^{\mathbf{M}}$. We name these restricted LSCM maps BD-LSCM (for Bounded Distortion). Figure 2 (c),(d) show deformation results with BD-LSCM with the conformal distortion bound set to $C = 1.5$. The blow-ups (b),(d) show specific areas where the conformal distortion of LSCM is high and alleviated by BD-LSCM. The initial frame field \mathbb{M} for $\mathcal{F}_C^{\mathbf{M},\mathbb{M},\infty}$ was extracted from the LSCM solution using eq. (4.21).

The color scheme (which is fixed for the rest of the paper, unless otherwise specified) should be interpreted as follows: redness indicates conformal distortion (**dark red** - high, **gray** - low), faces with **green** boundary means conformal distortion greater than 25 ($D(f_j) > 25$), **blue** boundary means conformal distortion greater than 50, and **yellow** boundary indicates flipped (also with **purple** interior) or degenerate triangles.

As-rigid-as-possible (ARAP) planar deformations. Similarly to ASAP, we can optimize the ARAP energy [Igarashi et al. 2005; Sorkine and Alexa 2007; Liu et al. 2008; Chao et al. 2010] over $\mathcal{F}_C^{\mathbf{M},\mathbb{M},\infty}$: given an initial deformation $\phi^* : \mathbf{M} \rightarrow \mathbb{R}^2$, we construct a frame field \mathbb{M} over \mathbf{M} using eq. (4.21). The ARAP energy is then $E_{\text{ARAP}}(\phi) = \sum_j [|\alpha_j - 1|^2 + |\beta_j|^2] \text{Area}(f_j)$. Indeed, since the affine transforms \mathcal{A}_j^* of ϕ^* , written in the frames \mathbb{M}_j , has $\alpha_j^* \in \mathbb{R}^+$ (we chose the frames that way), the above functional measures (for each face) deviation (in Frobenius norm) from the rigid map derived from \mathcal{A}_j^* by projecting its similarity part (α_j^*) on the planar rotation group (S^1) and throwing away its anti-similarity component (β_j^*); this is actually the closest rotation to \mathcal{A}_j^* .

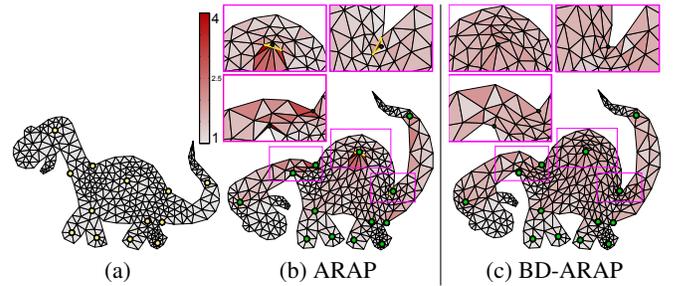


Figure 3: As-Rigid-As-Possible (ARAP) deformation of a dino model (shown in (a)) is shown in (b). The result of a Bounded distortion version of ARAP is in (c). Note that the ARAP energy actually decreased by a little more than 10% in (c) w.r.t. (b).

Figure 3(b) shows a deformation of a dino mesh (in (a)) by applying iterative ARAP algorithm similar to [Liu et al. 2008] (which provides a similar result to the other ARAP algorithms). ARAP deformations tend to concentrate distortion near constraints (especially interior ones), and flips and high conformal distortions are not uncommon (see blow-ups in (b)). In (c) we show the result of optimizing E_{ARAP} over $\mathcal{F}_C^{\mathbf{M},\mathbb{M},\infty}$ using the frame field \mathbb{M} extracted from the ARAP solution and setting $C = 2$. Note that the ARAP's high distortion spots and fold-overs that are now rectified. More surprising is the fact that the ARAP energy actually *decreased* by a little more than 10% in the BD-ARAP solution; this indicates that the ARAP algorithm got stuck in a local minimum.

Biharmonic planar warping. The Biharmonic energy was shown to be effective for creating smooth and intuitive warpings [Bookstein 1989; Jacobson et al. 2011]. In contrast to the ASAP and ARAP models, the Biharmonic warping is performed coordinate-wise and therefore does not have any control (also not on average) on the behavior of the differential and its distortion. For that reason, it is useful to restrict the Biharmonic energy to the bounded distortion spaces; for the Biharmonic energy, we adopt the linear finite element discretization of Jacobson *et al.* [2011]. That is, let L be the cotangent Laplacian matrix of mesh \mathbf{M} , and M the diagonal mass matrix. We also denote the vector $U = (u_1, u_2, \dots, u_{|V|})^t \in \mathbb{C}^n$. The Biharmonic energy is

$$E_{\text{biharm}}(\phi) = U^* L M^{-1} L U, \quad (6.1)$$

where superscript $*$ denotes the complex conjugate.

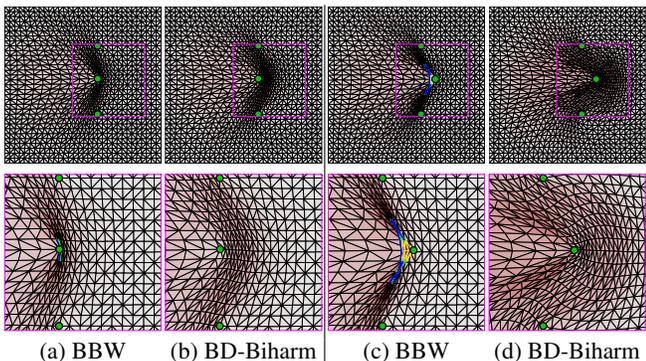


Figure 4: Warping a square domain using Jacobson et al. [2011] Bounded Biharmonic Weights (BBW) is shown in (a),(c). The Bounded Distortion (BD) Biharmonic solution is in (b),(d).

Figure 4 shows warpings of a square domain by fixing the square’s boundary and moving to the right one of the green points. In (a),(c) the results of the Bounded Biharmonic Weights (BBW) method [Jacobson et al. 2011] are shown, side by side with the Bounded-Distortion Biharmonic (BD-Biharm) solution in (b),(d), computed by optimizing (6.1) over $\mathcal{F}_C^{\mathbf{M}, \boxplus, \infty}$ with $C = 5$. The frame field \boxplus was initiated to be constant. Note that the BD-Biharmonic solution has better conformal distortion and avoids fold-overs. Lastly, note that Theorem 5.1 guarantees that the BD-Biharmonic map is a bijection of the square.

Dirichlet and stretch energy. Another popular convex quadratic energy used frequently in geometric processing is the Dirichlet energy that give rise to the discrete harmonic mappings [Pinkall and Polthier 1993; Eck et al. 1995; Floater and Hormann 2005]. Discrete harmonic mappings do not come with any guarantees regarding the conformal distortion they introduce, but for well shaped meshes Floater [2003] proved that they are injective if mapped to a convex domain. This is no longer true for non-convex target domains as shown in Figure 1(b). Optimizing the Dirichlet energy over $\mathcal{F}_C^{\mathbf{M}, \boxplus, \infty}$ guarantees both bounded conformal distortion and bijectivity even when mapping to non-convex and/or multiply-connected domains, see (d) in Figure 1 where we used $C = 5$ and set the initial frames from the discrete harmonic map. The Dirichlet energy can be expressed in our notations as $E_{\text{DIR}}(\phi) = \sum_j [|\alpha_j|^2 + |\beta_j|^2] \text{Area}(f_j)$, and in this example the energy increased by only 4% in the bounded distortion harmonic map version. This is not surprising as the discrete harmonic map is the global minimizer of the Dirichlet energy. Figure 1 (c) shows optimization of the L^1 stretch energy $E_{\text{stretch}_1}(\phi) =$

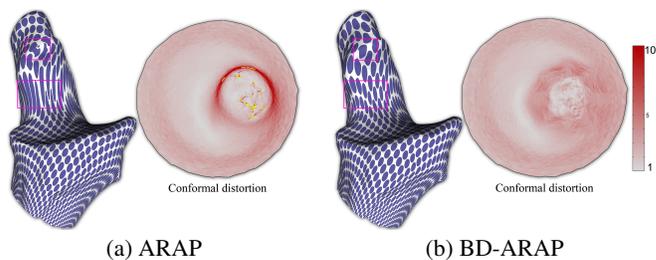


Figure 6: Parametrization of a metatarsal bone surface using Liu et al. [2008] ARAP method is shown in (a). Optimizing the same energy over the space of bounded distortion mappings (BD-ARAP), is shown in (b). BD-ARAP avoids the fold-overs and provides a guarantee on the maximal conformal distortion at the cost of slightly raising the ARAP minimal energy by a little less than 1.5%.

$\sum_j \Gamma_j \text{Area}(f_j) = \sum_j [|\alpha_j| + |\beta_j|] \text{Area}(f_j)$ ([Sander et al. 2001] introduced the L^∞ version of this energy) over the space $\mathcal{F}_C^{\mathbf{M}, \boxplus}$ (as this is L^1 energy). Note that the result has the same bounds as the BD-Harmonic solution but, as typical with L^1 norm minimizations in image processing, tend to concentrate distortion.

6.2 Surface parametrization

The algorithms we used for planar warping are readily adapted for surface parametrization. We now make the necessary adjustments and compare to the relevant state-of-art.

Conformal parametrization. Similarly to ASAP planar deformations, for conformal parameterizations we have optimized the LSCM energy over $\mathcal{F}_C^{\mathbf{M}, \boxplus, \infty}$. The frame field \boxplus was initialized via the gradient of a discrete harmonic-type function. Figure 5 compares our result with both Angle Based Flattening (ABF++) [Sheffer et al. 2005] and Conformal Equivalence of Triangle Meshes (CETM) [Springborn et al. 2008] on a collection of seven meshes. We set the conformal bound to be $C = 5$ (except in two cases where we had to increase it as no feasible solution was found). The color scheme is as described before, the conformal distortion color bars (for the left surface, and all the rest) are shown. For every mesh we show in parenthesis the L^1 (left) and maximum (right, in bold) norms of its conformal distortion. (for L^1 computation we excluded degenerate triangles in ABF++ and CETM.) BD-LSCM achieves significantly lower maximal conformal distortion, and usually lower L^1 distortion (for well-shaped meshes, like the left-most tooth surface, the L^1 is similar for all three methods). Note that a single LSCM parametrization result is shown in the ABF++ row (fandisk) as ABF++ did not find a feasible solution.

As-rigid-as-possible parametrization can be performed by repeating the procedure we introduced above for ARAP planar deformations. Figure 6 shows a comparison of Liu et al. [2008] ARAP parametrization with our result, achieved by optimization of the ARAP energy over $\mathcal{F}_C^{\mathbf{M}, \boxplus, \infty}$. Areas of high conformal distortion and flipped triangles in the ARAP result are marked. The BD-ARAP alleviates these issues, as before.

6.3 Controlling singular values/ area distortion

The bounded conformal distortion spaces $\mathcal{F}_C^{\mathbf{M}, \boxplus}$, $\mathcal{F}_C^{\mathbf{M}, \boxplus, \infty}$ can be further restricted to include also bounded area distortion guarantees. More generally, we describe here a way to construct uniform lower and upper bounds to the singular values of the affine maps over all faces. That is, given constants $0 \leq s \leq S$ we want the singular

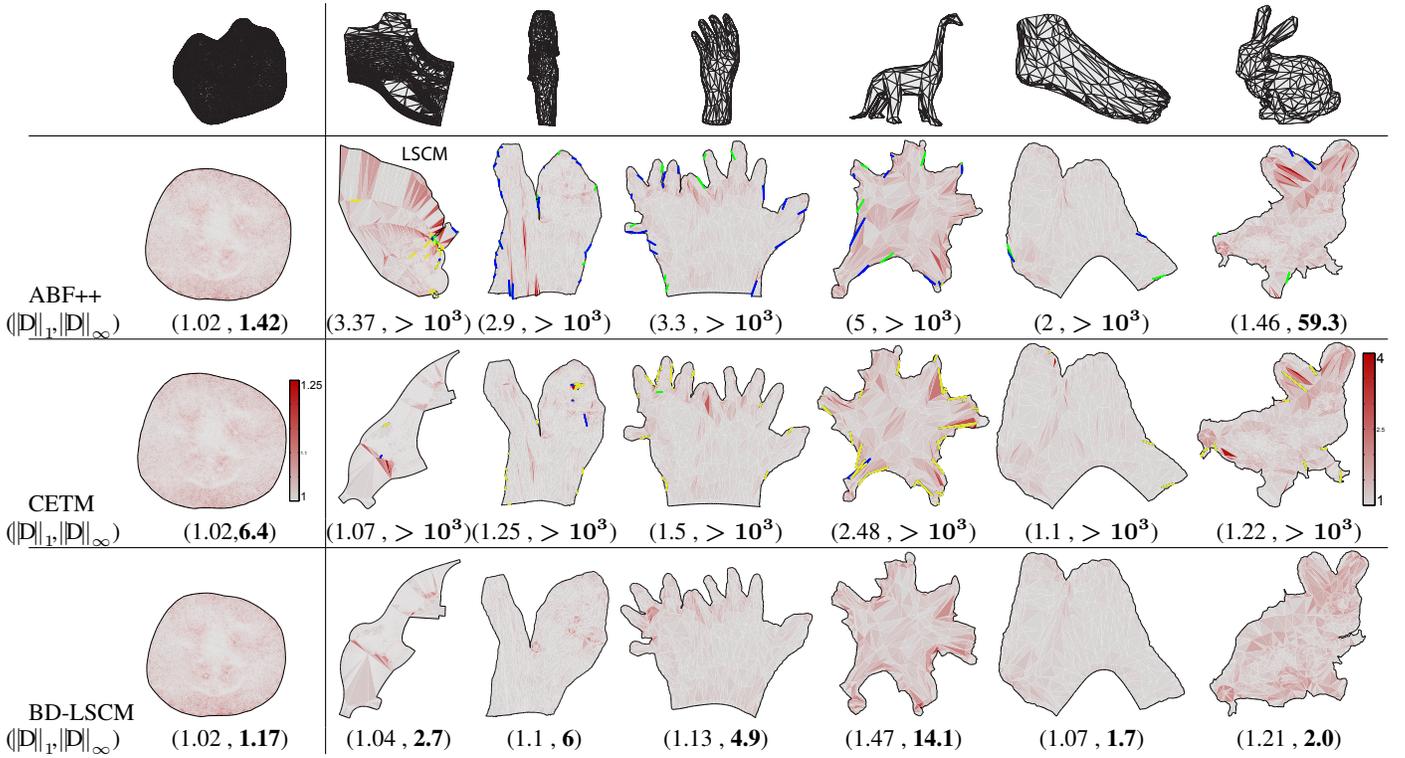


Figure 5: Parameterizations of seven surfaces (top row) cut to disk topology with Angle-Based Flattening (ABF++) [Sheffer et al. 2005], Conformal Equivalence of Triangle Meshes (CETM) [Springborn et al. 2008], and Bounded Conformal Distortion LSCM (BD-LSCM) are shown. We show the UV embeddings where faces are colored by their conformal distortion (green, blue, and yellow indicates high, very high and degenerate triangles, respectively), see text for detailed color scheme description. L^1 and maximal conformal distortion (bold) are shown in parenthesis below each flattening. ($> 10^3$ indicates maximal conformal distortion higher than 10^3 .) Both ABF++ and CETM perform well in terms of average conformal distortion, however few faces often sustain large conformal distortion. BD-LSCM has significantly lower bound on the maximal conformal distortion.

values γ_j, Γ_j of all affine map \mathcal{A}_j to satisfy

$$s \leq \gamma_j \leq \Gamma_j \leq S. \quad (6.2)$$

This can be achieved by adding the following two equations to eqs. (4.15)-(4.17):

$$|\alpha_j| \leq S - \frac{C-1}{C+1} r_j \quad (6.3)$$

$$r_j \geq \frac{C+1}{2} s. \quad (6.4)$$

Note that the first constraint is again a quadratic cone (and hence convex) and the second constraint is a simple linear inequality. Therefore, these constraints carve out a convex piece out of the convex spaces $\mathcal{F}_C^{\mathbf{M}, \mathbb{B}}, \mathcal{F}_C^{\mathbf{M}, \mathbb{B}, \infty}$. To see how (6.3), (6.4) imply eq. (6.2) note that (4.7), (4.15) and (6.3) imply $\Gamma_j = |\alpha_j| + |\beta_j| \leq \frac{C-1}{C+1} r_j + S - \frac{C-1}{C+1} r_j = S$. On the other hand, combining (4.8) with (4.15), (4.16), (4.17) and (6.4) gives $\gamma = |\alpha_j| - |\beta_j| \geq r_j \left(1 - \frac{C-1}{C+1}\right) = r_j \frac{2}{C+1} \geq s$. Two restrictions on s, S are: $S/s \geq C, S \geq r_j \geq s$ for all j .

Figure 7(c) shows a deformation of a bird mesh (shown in (a)) using BD-LSCM (with $C = 1.5$) which does not punish scaling. In (b) we optimized the LSCM energy on $\mathcal{F}_C^{\mathbf{M}, \mathbb{B}, \infty}$ with the addition of eqs. (6.4), (6.3) where we set $S = 1.25\sqrt{C}$, and $s = 1/S$. (The idea is that $S/s = 1.25^2 C$ which allows some freedom considering we restrict the conformal distortion to C .) The coloring in

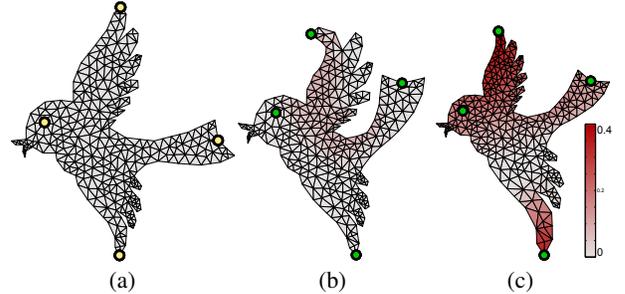


Figure 7: Deformation of a bird mesh (a) with (in (b)) and without (in (c)) the singular value bounds (6.3), (6.4). The coloring of each triangle is by the “rigidness” of its transform (As-Rigid-As-Possible energy) where redness indicates larger deviation from rigid transform.

this image is by the ARAP energy, namely each face is colored by $(\Gamma_j - 1)^2 + (\gamma_j - 1)^2$ which can be seen as the “rigidness” of the map on the faces. Note that the strict control over the singular values bounds the ARAP error per triangle.

6.4 Dependence on initial frame field

The main tools in this paper are the convex mapping spaces $\mathcal{F}_C^{\mathbf{M}, \mathbb{B}}, \mathcal{F}_C^{\mathbf{M}, \mathbb{B}, \infty}$, and therefore it is useful to understand their de-

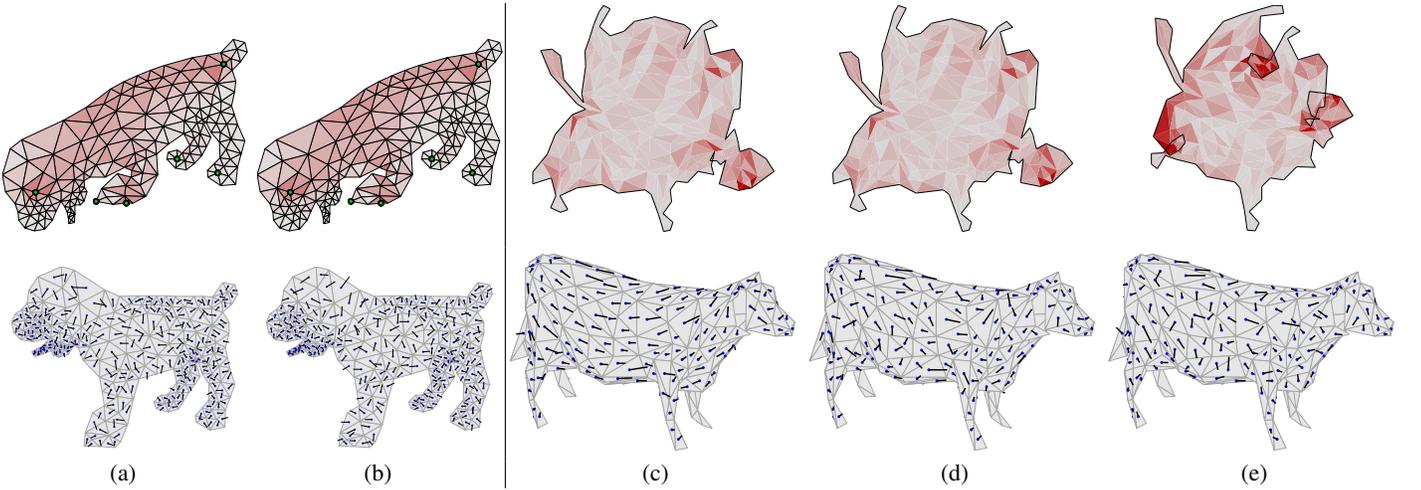


Figure 8: *Dependence of the mapping on the initial frame field. (a)-(b) show 2D deformations of a dog mesh (bottom) using BD-LSCM followed by BD-ARAP, with two initial frame fields (the 1st vectors of the frames \mathbb{T}_j , e_1^j , are shown in black on the original model in the bottom row). The resulting deformations are indistinguishable (top row). (c)-(e) show parameterizations of a cow model (bottom) with BD-LSCM. In (c) the initial frame field is set by the gradient field of a discrete harmonic-type function (bottom). In (d) the initial frame field of (c) is perturbed by random rotations of angles in the range $(-0.4\pi, 0.4\pi)$ and produces similar parameterization to (c). In (e) the discrete harmonic field is perturbed by random rotations in the range $(-0.7\pi, 0.7\pi)$. In this case the algorithm did not find a feasible solution.*

pendence on their unknown parameter, namely the frame field \mathbb{T} . On the theoretical side, Proposition 4.1 provides a characterization of the restriction on the rotation angle w.r.t. the initial rigid embedding of the faces defined by the frame-field. However, this analysis does not cover the case in which we allow resetting the frame-field and resolving (see the fourth step in the meta-algorithm description in the beginning of this section). In this subsection we present experimental results demonstrating the affect of perturbing the initial frame field on the optimized mappings.

When using planar deformations with positional constraints we found that allowing several iterations (we used four feasible iterations here) tend to lead to results independent of the initial frame field. Figure 8 (a),(b) show two deformations with the same positional constraints and bound on conformal distortion ($C = 2$) but with two randomly generated frame fields (depicted with black lines on the bottom row). Note that the final deformations on the top row are indistinguishable. However, in cases we did not incorporate positional constraints, such as free boundary parameterizations, we encountered stronger dependency on the initial frame field. Figure 8 (c)-(e) show three parameterizations of a low polygon count cow model where in (c) we used a discrete harmonic-type potential function to define the frame field, in (d) we perturbed this frame field by applying a random rotation in each face with angle range in $(-0.4\pi, 0.4\pi)$, and in (e) we randomly rotated in the range $(-0.7\pi, 0.7\pi)$. Note that while the parameterizations in (c) and (d) are very similar, in (e) the optimizer could not find a feasible solution.

6.5 Optimal quasi-conformal mappings

A natural question is: what is the minimal conformal distortion C for which the problem of mapping a mesh under some constraints is feasible? Or simply, what is the optimal quasi-conformal mapping?

We will consider a version of this problem where we start with a framed mesh $\mathbf{M} = (\mathbf{V}, \mathbf{E}, \mathbf{F}, \mathbb{T})$ and a set of positional constraints and we wish to solve:

$$C^* = \min \left\{ C \mid \mathcal{F}_C^{\mathbf{M}, \mathbb{T}} \neq \emptyset \right\}. \quad (6.5)$$

Algorithm 1: Optimal CPL quasi-conformal mapping.

Input: Framed Mesh $\mathbf{M} = (\mathbf{V}, \mathbf{E}, \mathbf{F}, \mathbb{T})$, $\mathbf{M} \subset \mathbb{R}^d$, $d = 2, 3$,
Tolerance $\varepsilon > 0$,
Point (or other) constraints

Output: extremal CPL-QC map $\phi^* \in \mathcal{F}_C^{\mathbf{M}, \mathbb{T}}$

set $C = C_0$ (e.g., 10), $C_{left} = 1$, $C_{right} = \infty$

while $C_{right} - C_{left} > \varepsilon$ **do**

Solve convex feasibility problem $\phi \in \mathcal{F}_C^{\mathbf{M}, \mathbb{T}}$

if *feasible* **then**

set $\phi^* = \phi \in \mathcal{F}_C^{\mathbf{M}, \mathbb{T}}$

forall faces $f_j //$ Field update

do

Update frame \mathbb{T}_j using ϕ^* , according to eq. (4.21)

$C_{right} = C$

else

$C_{left} = C$

if $C_{right} == \infty$ **then**

$C = 2C_{left}$

else

$C = \frac{1}{2} (C_{left} + C_{right})$

return ϕ^*

Then, any map $\phi^* \in \mathcal{F}_C^{\mathbf{M}, \mathbb{T}}$ would be optimal in that sense. In order to solve (6.5) we will use a simple modification of the bisection algorithm on the conformal bound C described in Algorithm 1. For finding a feasible $\phi \in \mathcal{F}_C^{\mathbf{M}, \mathbb{T}}$ we use MOSEK conic-programming [Andersen and Andersen 1999].

It is interesting to investigate the connection between the resulting map ϕ^* from Algorithm 1 and the (classical) extremal Quasi-Conformal (QC) mapping. Extremal QC mappings minimize the maximal conformal distortion in a given class of QC mappings and have fascinating mathematical theory and far reaching consequences [Imayoshi and Taniguchi 1992]. In practice, extremal QC map can simply be viewed as the most conformal solution to

a given mapping problem. The reason one would want to consider extremal QC maps in computer graphics is that "perfect" conformal maps have only a small number of degrees of freedom and cannot bijectively satisfy more than a few constraints simultaneously (see Section 2). We are not aware of any algorithm approximating extremal QC maps (for more than four interpolation points or a quad [Lipman et al. 2012]), and researchers have mainly focused on computing a QC map, not necessarily the optimal one, which is already a challenging problem [Gaidashev and Khmelev 2008; Zeng et al. 2009].

A first theoretical step is provided in the next theorem, describing a certain feasibility result:

Theorem 6.1. *Let $\Omega \subset \mathbb{C}$ be a planar domain with polygonal boundary. Let $\psi \in W_\infty^2(\Omega)$ be a QC map with maximal conformal distortion $D(\psi) = C^\psi$. We will also assume ψ has no critical points in Ω , that is $\gamma_\psi(z) \geq c > 0$ (a.e. in Ω). Let $\mathbf{M}^h = (\mathbf{V}^h, \mathbf{E}^h, \mathbf{F}^h)$ be a non-degenerate triangular mesh² subdividing Ω with maximal face diameter of $h > 0$.*

Then, there exists a CPL map $\phi^+ \in \mathcal{F}^{\mathbf{M}}$ defined over mesh \mathbf{M} with maximal conformal distortion $D(\phi^+) = C^\psi + O(h)$. In other words, $\mathcal{F}_{C^\psi + O(h)}^{\mathbf{M}} \neq \emptyset$.

Intuitively, this theorem implies that if the extremal map $\psi : \Omega \rightarrow \mathbb{C}$ is regular enough (belongs to the Sobolev space $W_\infty^2(\Omega)$), and the mesh \mathbf{M}^h of the domain Ω is fine enough, then there is CPL map ϕ^+ (which is simply the CPL interpolant of ψ over \mathbf{M}) that has conformal distortion $O(h)$ away from the optimal conformal distortion of ψ . This result implies that we can find a CPL map with maximal conformal distortion $O(h)$ close to optimal if we could search globally in the entire space of bounded distortion CPL maps, but we don't know how (this is a non-convex space). Instead we search in $\mathcal{F}_C^{\mathbf{M}, \mathbb{B}}$, and so we need to make sure our frame field \mathbb{B} is chosen *close enough* to ensure convergence. We cannot guarantee that as we take our frame field from some initial guess (in this case we extracted it from the discrete harmonic mapping). In a sense we find optimal solution in a certain known proximity to our initial guess (the proximity is understood from Proposition 4.1). This is the reason we have in Algorithm 1 a field update stage.

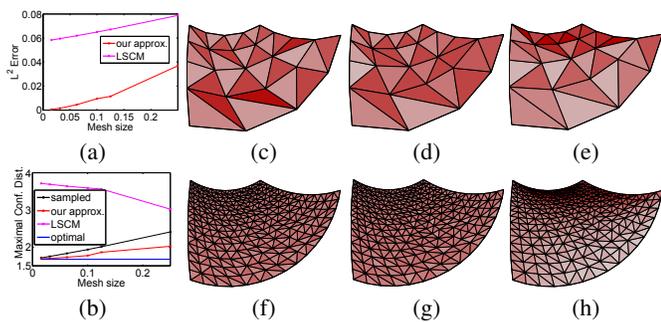


Figure 9: Convergence of optimal QC to the extremal QC map. (a) shows L^2 convergence plot of the approximate map to the reference map (red curve). The LSCM error is shown in purple. (b) shows a plot of the maximal conformal distortion of the optimal smooth map (blue), its sampled CPL version over the mesh (black), our approximation (red), and LSCM (purple); (c-e) and (f-h) show two different mesh resolution (from left to right): the sampled analytic extremal QC map, our approximation, and LSCM.

However, in practical scenarios we have considered, evidence of convergence to the extremal map was observed, as described next.

²A non-degenerate mesh has uniform bound on the faces' chunkiness parameter (see Definition 6.2.16 in [Brenner and Scott 2008])

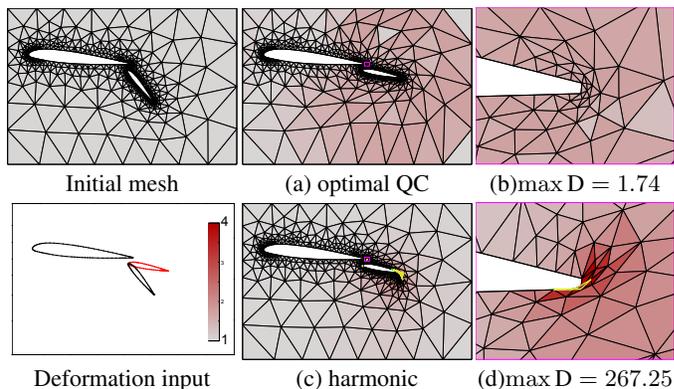


Figure 10: Deformation of multiply-connected airfoil domain with discrete harmonic map (c),(d) and optimal quasi-conformal map (a),(b).

Figure 9 shows a convergence plot (a)(red curve) of the mappings produce by Algorithm 1. Our approximations are computed over a series of meshes of a square with decreasing mesh size and the L^2 deviation from an analytical solution are shown. The analytical solution (of deforming a square by moving two of its corners) is taken from [Lipman et al. 2012] where a formula for extremal QC mappings is provided for the special case of interpolating four points in the plane with certain (conformal-periodic) boundary conditions. The boundary vertices in our approximation are constrained according to the analytic solution (we simply set the relevant positional constraints in $\mathcal{F}_C^{\mathbf{M}, \mathbb{B}}$). For comparison we solve Least-Squares Conformal Maps with the same boundary conditions and show its L^2 error curve as well (in purple). The rows (c)-(e) and (f)-(h) show the results for two different mesh resolutions, where (c),(f) show the analytic solution interpolated over the mesh, (d),(g) our approximation resulted from Algorithm 1, and (e),(h) the LSCM result. Note that both the plot and the meshes suggest convergence of our approximation scheme. In (b) we show the maximal conformal distortion as a function of the mesh size. The optimal solution's conformal distortion is constant (in blue), while the sampled (interpolated) optimal solution has linear convergence in the mesh size (in black) as Theorem 6.1 implies. As expected (but not proved!) our solution's curve is bounded below the curve of the sampled (interpolated) optimal solution and therefore has also empirical linear convergence to the optimal distortion.

Although rather costly to compute (we discuss timing at the end of this section) the optimal maps are of high quality (in terms of conformal distortion), and we believe that they will find their applications where the low distortion of the map is more important than the speed of its computation. For example, Figure 10 (a),(b) show the result of Algorithm 1 in computing a map between two multiply-connected planar domains of an airfoil used for Finite-Element analysis. (c),(d) show the discrete harmonic mapping which suffers from fold-overs and high conformal distortion. The optimal QC map has low (relative to the discrete harmonic map) maximal conformal distortion (1.74) and it is guaranteed to be bijective (by Theorem 5.2). In this case we used a constant frame field $\{1, \mathbf{i}\}$ since the frame field of the discrete harmonic initial guess was very distorted in some areas (here we used MOSEK rather than Matlab's QP for optimization).

6.6 2D mesh improvement

The goal of planar meshing algorithms is to produce triangulations of polygonal domains with guarantees on the shape of the triangles. Remeshing of 2D planar domain can be cast into planar mapping problem as follows: given a mesh \mathbf{M} of a planar domain $\Omega \subset \mathbb{R}^2$,

we can think of it as an image under some map ϕ of a hypothetical mesh \mathbf{M}^* with the same connectivity as \mathbf{M} but perfect equilateral faces. We don't really care if \mathbf{M}^* can or cannot be embedded in a specific Euclidean space; in fact, we take the Riemannian point of view - we change the metric of \mathbf{M} to define \mathbf{M}^* , that is we choose a frame field \mathbb{T} on \mathbf{M} such that representing each face f_j in its local frame gives an equilateral which is closest to the actual f_j . We compute such frames by simply fitting, in the least-squares sense, equilateral to each face. Then, we use Algorithm 1 to find an optimal QC map $\phi^* : \mathbf{M}^* \rightarrow \Omega$ with as small as possible conformal distortion C^* . The image $\tilde{\mathbf{M}} = \phi^*(\mathbf{M}^*)$ is our optimized mesh where we know that each face was distorted by no more than C^* from a perfect equilateral, and ϕ^* is a bijection.

Two comments are: we set the boundary positional constraints by requiring that $\phi^*(\partial\mathbf{M}^*) = \partial\mathbf{M}$ (i.e., the boundary of the mesh \mathbf{M} is fixed). Second, we note that all the equations defining $\mathcal{F}_C^{\mathbf{M}, \mathbb{T}}$ can be applied to the hypothetical mesh \mathbf{M}^* , nowhere did we require an actual embedding to define $\mathcal{F}_C^{\mathbf{M}, \mathbb{T}}$.

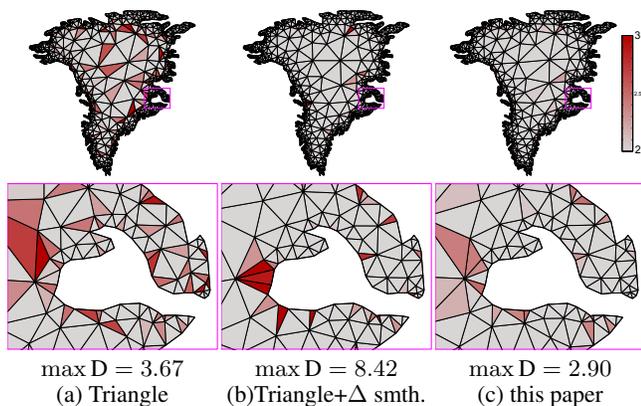


Figure 11: 2D meshing of Greenland's shoreline polygon: Shewchuk's Triangle [1996] result is shown in (a), applying further Laplacian smoothing results in (b). In (c) we optimize the Triangle's mesh (in (a)) using the optimal QC algorithm 1.

Figure 11 shows meshing of Greenland's shoreline polygon. We show meshing using the Triangle software [Shewchuk 1996] in (a), and the Laplacian smoothing of the Triangle's mesh in (b). In (c), we show our optimization applied to the mesh created with Triangle. We color the faces based on conformal distortion w.r.t. an equilateral. Although the Laplacian smoothing regularized the mesh it increased the maximal conformal distortion; our optimization regularized the mesh as well as reduced the maximal conformal distortion to 2.9 (from 3.67 of the initial Triangle's mesh in (a)). The histogram of the individual faces' conformal distortion for the three results is shown in the inset. Note that the histogram of our remeshing (in blue) ends slightly before hitting $D = 3$.

Timing. Most of our code is written in Matlab and is not optimized for speed. We ran our experiments on a 3.4GHz PC processor. For optimizing over $\mathcal{F}_C^{\mathbf{M}, \mathbb{T}}$ (L^1 Stretch, QC optimal, and remeshing) we used external calls to MOSEK [Andersen and Andersen 1999], and for all the rest we had quadratic energies that were optimized over $\mathcal{F}_C^{\mathbf{M}, \mathbb{T}, \infty}$ with Matlab's Quadratic Program-

ming. Typical optimization times are: BD-ARAP (Figure 3) 1.1sec for $|V| = 301$; BD-LSCM (Figure 5) 2.86sec for $|V| = 1056$, and 37.6sec for $|V| = 8600$; BD-ARAP (Figure 6) 41sec for $|V| = 2530$. In Figure 1, $|V| = 2193$ and MOSEK optimization of the L^1 Stretch took 8.2sec, while Matlab's optimization of Dirichlet took on this example 20.6sec. Optimal QC computation are more expensive due to numerous applications of the optimization by Algorithm 1: the map in Figure 10 took 17sec for $|V| = 600$, and the mesh optimization in Figure 11 took 14sec for $|V| = 1200$.

7 Limitations

There are two main limitations to our technique. The first, and more obvious one, is the computational complexity that stems from the non-linear formulation of the bounded distortion mapping spaces.

The second is more theoretical: we did not address the question of feasibility in its general form. Theorem 6.1 provides a first step in this direction, but it still does not answer the question: what are the conditions that assure that for a given mesh and a set of mapping constraints there exists a solution $\mathcal{F}_C^{\mathbf{M}, \mathbb{T}} \neq \emptyset$ (or $\mathcal{F}_C^{\mathbf{M}, \mathbb{T}, \infty} \neq \emptyset$). For example, in Figure 5 we could not find a feasible solution for the dino mesh with $C = 5$, and needed to raise the bound to $C = 15$. Answering this question will provide a better picture of how much flexibility we have in mapping triangular meshes.

8 Conclusions and future work

This paper introduced convex mapping spaces for triangular meshes with guarantees. Our focus was controlling worst-case conformal distortion in various existing mapping algorithms. In addition we examined applications to bijective mappings, extremal quasi-conformal mappings, controlling the singular values, and planar remeshing.

The bounded distortion mappings introduced in this paper satisfies the *continuous* definition of quasi-conformal mappings (since QC maps allow weaker differentiability than, say, standard conformal mappings). This provides the opportunity of using the well developed continuous QC theory directly on meshes, a fact which opens many future research directions. We also plan to search for more applications in geometry processing for these spaces. One obvious direction is to build bounded distortion intrinsic mappings between two surfaces embedded in 3D, a core element in many shape analysis applications.

Acknowledgements

The author would like to thank Denis Zorin for useful discussions, Alec Jacobson and Olga Sorkine for the implementation of BBW and ARAP, Boris Springborn for the implementation of CETM, and the anonymous reviewers for their useful comments and suggestions. Most 3D Models taken from <http://alice.loria.fr/index.php/software/7-data/37-unwrapped-meshes.html>.

References

- AHLFORS, L. 1966. *Lectures on quasiconformal mappings*. University lecture series. American Mathematical Society.
- ALEXA, M., COHEN-OR, D., AND LEVIN, D. 2000. As-rigid-as-possible shape interpolation. In *Proceedings of SIGGRAPH '00*, ACM, New York, NY, USA, 157–164.
- AMENTA, A. B., BERN, M. W., AND EPPSTEIN, D. 1997. Optimal point placement for mesh smoothing. In *Proc. 8th Symp. Discrete Algorithms*, ACM and SIAM, 528–537.
- ANDERSEN, E. D., AND ANDERSEN, K. D. 1999. *The MOSEK*

- interior point optimization for linear programming: an implementation of the homogeneous algorithm.* Kluwer Academic Publishers, 197–232.
- ARONOV, B., SEIDEL, R., AND SOUVAINE, D. 1993. On compatible triangulations of simple polygons. *Comput. Geom. Theory Appl.* 3 (June), 27–35.
- BEN-CHEN, M., GOTSMAN, C., AND BUNIN, G. 2008. Conformal flattening by curvature prescription and metric scaling. *Computer Graphics Forum* 27.
- BERN, M. W., AND EPPSTEIN, D. 1995. Mesh generation and optimal triangulation. In *Computing in Euclidean Geometry*, D.-Z. Du and F. K.-M. Hwang, Eds., second ed., no. 4 in Lecture Notes Series on Computing. World Scientific, 47–123.
- BOOKSTEIN, F. L. 1989. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell.* 11 (June), 567–585.
- BRENNER, S., AND SCOTT, L. 2008. *The mathematical theory of finite element methods.* Texts in applied mathematics. Springer.
- CHAO, I., PINKALL, U., SANAN, P., AND SCHRÖDER, P. 2010. A simple geometric model for elastic deformations. *ACM Trans. Graph.* 29 (July), 38:1–38:6.
- DESBRUN, M., MEYER, M., AND ALLIEZ, P. 2002. Intrinsic Parameterizations of Surface Meshes. *Computer Graphics Forum* 21.
- ECK, M., DE ROSE, T., DUCHAMP, T., HOPPE, H., LOUNSBERY, M., AND STUETZLE, W. 1995. Multiresolution analysis of arbitrary meshes. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, SIGGRAPH '95, 173–182.
- FLOATER, M. S., AND HORMANN, K. 2005. Surface Parameterization: a Tutorial and Survey. *Advances in Multiresolution for Geometric Modelling*, 157–186.
- FLOATER, M. S. 2003. One-to-one piecewise linear mappings over triangulations. *Mathematics of Computation* 72, 685–696.
- GAIDASHEV, D., AND KHMELEV, D. 2008. On numerical algorithms for the solution of a beltrami equation. *SIAM J. Numer. Anal.* 46 (May), 2238–2253.
- GORTLER, S. J., GOTSMAN, C., AND THURSTON, D. 2006. Discrete one-forms on meshes and applications to 3d mesh parameterization. *Comput. Aided Geom. Des.* 23, 2 (Feb.).
- HORMANN, K., AND GREINER, G. 2000. MIPS: An efficient global parameterization method. In *Curve and Surface Design: Saint-Malo 1999*, P.-J. Laurent, P. Sablonnière, and L. L. Schumaker, Eds., Innovations in Applied Mathematics. Vanderbilt University Press, Nashville, TN, 153–162.
- IGARASHI, T., MOSCOVICH, T., AND HUGHES, F. J. 2005. As-rigid-as-possible shape manipulation. *ACM Trans. Graph.* 24, 1134–1141.
- IMAYOSHI, Y., AND TANIGUCHI, M. 1992. *An introduction to Teichmüller spaces.* Springer-Verlag.
- JACOBSON, A., BARAN, I., POPOVIĆ, J., AND SORKINE, O. 2011. Bounded biharmonic weights for real-time deformation. *ACM Transactions on Graphics (proceedings of ACM SIGGRAPH)* 30, 4, 78:1–78:8.
- JIN, M., KIM, J., AND GU, X. D. 2007. Discrete surface ricci flow: Theory and applications. In *In IMA Conference on the Mathematics of Surfaces*, 209–232.
- JOHNEN, A., REMACLE, J.-F., AND GEUZAIN, C. 2012. Geometrical validity of curvilinear finite elements. In *Proceedings of the 20th International Meshing Roundtable*. 255–271.
- KHAREVYCH, L., SPRINGBORN, B., AND SCHRÖDER, P. 2006. Discrete conformal mappings via circle patterns. *ACM Trans. Graph.* 25 (April), 412–438.
- LEVY, B., PETITJEAN, S., RAY, N., AND MAILLO T, J. 2002. Least squares conformal maps for automatic texture atlas generation. In *ACM SIGGRAPH conference proceedings*, ACM, Ed.
- LIPMAN, Y., LEVIN, D., AND COHEN-OR, D. 2008. Green coordinates. *ACM Trans. Graph.* 27, 3.
- LIPMAN, Y., KIM, V., AND FUNKHOUSER, T. 2012. Simple formulas for quasiconformal plane deformations. *ACM Transactions on Graphics, accepted for publication.*
- LIU, L., ZHANG, L., XU, Y., GOTSMAN, C., AND GORTLER, S. J. 2008. A local/global approach to mesh parameterization. In *Proceedings of the Symposium on Geometry Processing*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, SGP '08, 1495–1504.
- MAILLOT, J., YAHIA, H., AND VERROUST, A. 1993. Interactive texture mapping. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, ACM, New York, NY, USA, 27–34.
- PINKALL, U., AND POLTHIER, K. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2, 15–36.
- RUPPERT, J. 1995. A delaunay refinement algorithm for quality 2-dimensional mesh generation. In *Selected papers from the fourth annual ACM SIAM symposium on Discrete algorithms*, Academic Press, Inc., Orlando, FL, USA, 548–585.
- SANDER, P. V., SNYDER, J., GORTLER, S. J., AND HOPPE, H. 2001. Texture mapping progressive meshes. In *Proceeding of SIGGRAPH '01*, ACM, New York, NY, USA, 409–416.
- SCHAEFER, S., MCPHAIL, T., AND WARREN, J. 2006. Image deformation using moving least squares. *ACM Trans. Graph.* 25 (July), 533–540.
- SHEFFER, A., LÉVY, B., MOGILNITSKY, M., AND BOGOMYAKOV, A. 2005. Abf++: fast and robust angle based flattening. *ACM Trans. Graph.* 24 (April), 311–330.
- SHEFFER, A., PRAUN, E., AND ROSE, K. 2006. Mesh parameterization methods and their applications. *Found. Trends. Comput. Graph. Vis.* 2 (January), 105–171.
- SHEWCHUK, J. R. 1996. Triangle: Engineering a 2D Quality Mesh Generator and Delaunay Triangulator. In *Applied Computational Geometry: Towards Geometric Engineering*, M. C. Lin and D. Manocha, Eds., vol. 1148 of *Lecture Notes in Computer Science*. Springer-Verlag, May, 203–222.
- SOLOMON, J., BEN-CHEN, M., BUTSCHER, A., AND GUIBAS, L. 2011. As-Killing-As-Possible Vector Fields for Planar Deformation. *Computer Graphics Forum* 30, 5, 1543–1552.
- SORKINE, O., AND ALEXA, M. 2007. As-rigid-as-possible surface modeling. In *Proceedings of the fifth Eurographics symposium on Geometry processing*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 109–116.
- SORKINE, O., COHEN-OR, D., GOLDENTHAL, R., AND LISCHINSKI, D. 2002. Bounded-distortion piecewise mesh parameterization. In *Proceedings of IEEE Visualization*, IEEE Computer Society, 355–362.
- SPRINGBORN, B., SCHRÖDER, P., AND PINKALL, U. 2008. Conformal equivalence of triangle meshes. *ACM Trans. Graph.* 27 (August), 77:1–77:11.
- STEPHENSON, K. 2005. *Introduction to circle packing: the theory of discrete analytic functions.* Cambridge University Press.

TUTTE, W. T. 1963. How to Draw a Graph. *Proceedings of the London Mathematical Society* s3-13, 1, 743–767.

WEBER, O., AND GOTSMAN, C. 2010. Controllable conformal maps for shape deformation and interpolation. *ACM Trans. Graph.* 29, 78:1–78:11.

ZENG, W., AND GU, X. D. 2011. Registration for 3d surfaces with large deformations using quasi-conformal curvature flow. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, 2457–2464.

ZENG, W., LUO, F., YAU, S. T., AND GU, X. D. 2009. Surface quasi-conformal mapping by solving beltrami equations. In *Proceedings of the 13th IMA International Conference on Mathematics of Surfaces XIII*, Springer-Verlag, Berlin, Heidelberg, 391–408.

Appendix A

In this appendix we prove Theorem 5.2 which is more general and includes Theorem 5.1. We start with three Lemmas.

Lemma A.1. (*local bijectivity*) Any map $\phi \in \mathcal{F}_+^{\mathbf{M}}$ is bijective on all its quads (a quad is the union of two adjacent faces sharing an edge).

Proof. Let $f_j, f_k \in \mathbf{F}$ be two faces sharing an edge $e_\ell \in \mathbf{E}$. Since their frames \boxplus_j, \boxplus_k are positively oriented (i.e., consistent with the orientation of the mesh) and orthonormal, they define a congruent, positive oriented embedding of each of these faces in the plane $\Delta_j = \Delta([v_{j_1}], [v_{j_2}], [v_{j_3}])$, $\Delta_k = \Delta([v_{k_1}], [v_{k_2}], [v_{k_3}]) \subset \mathbb{R}^2$ (respectively). The affine maps $\mathcal{A}_j, \mathcal{A}_k$ are, by definition of $\mathcal{F}_+^{\mathbf{M}}$, orientation preserving and coincide on the common edge e_ℓ . This implies that $\mathcal{A}_j(\Delta_j), \mathcal{A}_k(\Delta_k)$ lie on different sides of $\mathcal{A}_j(e_\ell) = \mathcal{A}_k(e_\ell)$. \square

Lemma A.2. (*interior vertex*) Any map $\phi \in \mathcal{F}_+^{\mathbf{M}}$ takes some neighborhood of any interior vertex $v \in \mathbf{V}$ to a neighborhood of its image $\phi(v)$.

Proof. Let $v \in \mathbf{V}$ be an interior vertex (i.e., not on the boundary). Let e_1, e_2, \dots, e_n be the ordered (by the orientation) edges that touch v . The angle $\theta_{k, k+1}$ between $\phi(e_k)$ and $\phi(e_{k+1})$ is in $(0, \pi)$ (otherwise the map would be orientation reversing or degenerate). If we sum the angles then $\sum_{\ell=1}^n \theta_{\ell, \ell+1} = 2m\pi$ by continuity of the map ϕ . Since all $\theta_{k, k+1} > 0$ we have $m > 0$ and so $m \in \{1, 2, 3, \dots\}$ and so the lemma is proved. \square

Lemma A.3. (*interior point*) Any map $\phi \in \mathcal{F}_+^{\mathbf{M}}$, takes some neighborhood of any interior point $p \in \mathbf{M}^\circ$ (superscript $^\circ$ denotes the interior of a set) to a neighborhood of its image $\phi(p)$.

Proof. The point p is either: 1) an interior vertex, 2) on an interior edge but not a vertex, or 3) in the interior of a face. Lemma A.2 takes care of case (1). Lemma A.1 takes care of case (2) (we use the fact that the inverse is a continuous map). Lastly, case (3) is proved since by construction $\mathcal{F}_+^{\mathbf{M}}$ is a non-degenerate affine map over each face. \square

Proof. (of **Theorem 5.2**) Denote by $\mathfrak{w}(w, P'_\ell)$ the winding number of P'_ℓ w.r.t. the point $w \notin \partial\Omega$. Given a point $w \notin \partial\Omega$ one can count the number of pre-images of w , namely $\text{Card}\phi^{-1}(w) = \text{Card}\{p \in \mathbf{M} \mid \phi(p) = w\}$ ³ using a version of the discrete argument principle (see for example [Stephenson 2005], Lemma 11.3, page 141). Indeed, let w be a point that does not intersect the image under ϕ of any edge, that is $w \notin \phi(E)$, where we denote here the

edge set $E = \cup_k e_k$. Since each face is mapped with an orientation preserving affine map, we have that

$$\text{Card}\phi^{-1}(w) = \sum_{j=1}^{|\mathbf{F}|} \mathfrak{w}(w, \partial\phi(f_j)), \quad (\text{A.1})$$

where $\partial\phi(f_j)$ denotes the oriented boundary curve of the image of face f_j . For every two neighboring faces $\phi(f_j), \phi(f_{j'})$ the winding number integral cancels on their common edge, and therefore

$$\sum_{j=1}^{|\mathbf{F}|} \mathfrak{w}(w, \partial\phi(f_j)) = \sum_{\ell=0}^L \mathfrak{w}(w, P'_\ell). \quad (\text{A.2})$$

Combining eq. (A.1) and eq. (A.2) gives the discrete argument principle:

$$\text{Card}\phi^{-1}(w) = \sum_{\ell=0}^L \mathfrak{w}(w, P'_\ell). \quad (\text{A.3})$$

We consider three cases: 1) w is in the unbounded component of $\Omega^c = \mathbb{C} \setminus \Omega$; 2) w is in one of the bounded components of Ω^c (i.e., holes); and 3) $w \in \Omega^\circ$. In the first case $\mathfrak{w}(w, \phi(P'_\ell)) = 0$ for all ℓ , while in the second case $\mathfrak{w}(w, \phi(P'_0)) = 1$ and there exists a unique $\ell \in \{1, 2, \dots, L\}$ such that $\mathfrak{w}(w, \phi(P'_\ell)) = -1$. Using the discrete argument principle eq. (A.3), for both cases, $\text{Card}\phi^{-1}(w) = 0$. Similar reasoning shows that for the third case $\text{Card}\phi^{-1}(w) = 1$.

Next, we show that ϕ maps the interior of \mathbf{M} to the interior of Ω . Otherwise we have an interior point $o \in \mathbf{M}^\circ$ mapped to the boundary or the complement of the domain, $\phi(o) \in \partial\Omega \cup \Omega^c$, and Lemma A.3 indicates the existence of a neighborhood of $\phi(o)$ in the image set $\phi(\mathbf{M})$. But, that contradicts the fact that for any $w \in \Omega^c \setminus \phi(E)$, $\text{Card}\phi^{-1}(w) = 0$. So $\phi(\mathbf{M}) \subset \Omega$. From the discrete argument principle above we also got $\Omega \setminus \phi(E) \subset \phi(\mathbf{M})$, and taking the closure of both sides we get $\Omega \subset \phi(\mathbf{M})$. So $\phi(\mathbf{M}) = \Omega$ (onto).

To prove injectivity assume (distinct) $p, p' \in \mathbf{M}$ such that $\phi(p) = \phi(p')$. p, p' cannot both belong to the boundary as the boundary is mapped bijectively. Since we showed ϕ maps the interior of \mathbf{M} to interior of Ω , the only remaining option is that both $p, p' \in \mathbf{M}^\circ$. If $\phi(p) = \phi(p') \in \phi(E)$ we can use Lemma A.3 to perturb p, p' a bit such that still $\phi(p) = \phi(p')$ but $\phi(p) \notin \phi(E)$. However we already showed that any $w \in \Omega \setminus \phi(E)$ has exactly one pre-image in \mathbf{M} , which leads to a contradiction. \square

Appendix B

We prove Proposition 4.1:

Proof. (of **Proposition 4.1**) Let $\mathcal{A}(z) = \alpha z + \beta \bar{z} + \delta$ be an affine map with conformal distortion $1 \leq c \leq C$, then from eq. (4.9) and orientation preserving property we get (by canceling $|\alpha|$ from numerator and denominator) $|\frac{\beta}{\alpha}| \leq \frac{c-1}{c+1}$. Now a simple calculation shows

$$\text{Re}(\alpha) = |\alpha| \cos(\arg(\alpha)) \geq |\alpha| \frac{C+1}{C-1} \frac{c-1}{c+1} \geq |\beta| \frac{C+1}{C-1},$$

where we used $|\arg(\alpha_j)| \leq \cos^{-1}\left(\frac{C+1}{C-1} \frac{c-1}{c+1}\right)$ in the second inequality, and $|\beta| \leq \frac{c-1}{c+1} |\alpha|$ in the last inequality. Take $r = \text{Re}(\alpha)$ and the proposition follows. \square

Lemma B.1. The L_∞ cone $|z|_\infty \leq R$, where $z \in \mathbb{C}$, is contained in the quadratic cone $|z| \leq \sqrt{2}R$, and contains the quadratic code $|z| \leq R$.

Proof. This is basically the equivalence relations of the infinity and 2-norm: let $z = x + iy$, $x, y \in \mathbb{R}$, then $|z| = \sqrt{x^2 + y^2} \leq \max\{|x|, |y|\} \sqrt{2} = |z|_\infty \sqrt{2}$, and $|z|_\infty \leq |z|$. \square

Proof. (of **Proposition 4.2**) This is a direct consequence of Lemma B.1. \square

³CardA denotes the number of elements in the set A.